

A Theory of Probabilistic Boosting, Decision Trees and Matryoshki

Etienne Grossmann
visiting at Université de Montréal
<etienne@isr.ist.utl.pt>

Abstract

We present a theory of boosting probabilistic classifiers. We place ourselves in the situation of a user who only provides a stopping parameter and a probabilistic weak learner/classifier and compare three types of boosting algorithms: probabilistic Adaboost, decision tree, and tree of trees of ... of trees, which we call *matryoshka*. “Nested tree,” “embedded tree” and “recursive tree” are also appropriate names for this algorithm, which is one of our contributions. Our other contribution is the theoretical analysis of the algorithms, in which we give training error bounds. This analysis suggests that the *matryoshka* leverages probabilistic weak classifiers more efficiently than simple decision trees.

1 Introduction

Ensembles of classifiers are a popular way to build a strong classifier by leveraging simple decision rules -weak classifiers. Many ensemble architectures have been proposed, such as neural networks, decision trees, Adaboost [1], bagged classifiers [2], random forests [3], trees holding a boosted classifier at each node [4], boosted decision trees... One drawback of ensemble methods is that they are often dispendious about the computational cost of the resulting classifier. For example, Adaboost [1], bagging [2], random forests [3] all multiply the runtime complexity, by a factor approximately proportional to the training time. This is not acceptable in applications involving large amounts of data and requiring low-complexity method, such as video analysis and data mining.

Many approaches have been proposed to deal with such situations. The cascade architecture, i.e. a degenerate decision tree, has become very popular [5] and has been intensely studied [6, 7]. However, cascades are mostly appropriate to detect rare exemplars of interest amongst a huge majority of uninteresting ones.

Decision trees, on the other hand, are better adapted to the case of balanced target classes. This advantage comes from their greater facility to decompose the input space into more manageable and useful subsets. In addition, their run-time complexity is approximately proportional to the logarithm of the training time. Counterbalancing these advantages, is the fact that decision trees tend to overfit the training data.

There exist many proposed methods to improve overfitting, for example pruning and smoothing, but the main recognized cause remains: data elements are passed to one only of the descendant of each node, whether during training, or at run-time. At run-time, one proposed solution is to pass examples along more than one child node [8].




Figure 1: Left: Construction of a matryoshka decision trees. Here, each tree has just two nodes, but other numbers are possible. Right: notation use to specify a path in a decision tree.

During training, it has been proposed [4] to pass down to all descendants the exemplars that lay within a fixed distance of the separating surface.

Our approach to avoid the hard split at each node is to consider that the examples have a certain probability -not necessarily always 0 or 1- of being passed to any descendant of the tree. That is, we study probabilistic decision trees [9], but pursue a different analysis from these last works. First, we show that probabilistic decision trees are eminently tractable within the framework of boosting¹.

We bound the expected misclassification error as a function of the number of nodes in the tree, in Section 4. This bound is very high when the probabilistic weak classifiers are very weak. Moreover, we present arguments that suggest that any bound using the same probabilistic weak learner hypothesis will necessarily be high.

However, we also note that the bound achieved with stronger weak classifiers is much better. In an attempt to strengthen our weak classifiers, we explore the possibility of assembling decision trees consisting of decision trees, the inner and the outer trees being built by the same algorithm. This is not the first time this idea is suggested, but we believe we are the first to show the theoretical benefits of doing so.

Continuing on the idea of embedding (or nesting) decision trees one into another, we propose to assemble trees of trees of ... of trees of probabilistic weak classifiers. This is similar to matryoshka dolls, with the difference that each tree contains more than one tree, rather than a single other doll. Figure 1, left, illustrates this concept. Our main contribution (Section 5.2) is to prove a greatly improved bound, reached by trees with exactly two nodes, each node being a tree with two nodes, and so on until the last nesting level, which holds two probabilistic weak classifiers.

Another merit of our study is that it proposes a methodology that is essentially parameterless. The user only needs to provide a probabilistic weak learner and a stopping criterion, such as the number of nodes or the error on the training dataset. If a stopwatch², is available during training, then we propose ways of using it. The freedom of parameters results partly from applying a principle of greedy error minimization.

Before presenting our study on decision trees, we define, in Section 2, the probabilistic weak learners that are the basis of this work. We then present, in Section 3,

¹The analogy between deterministic decision trees and boosting has been studied in [10].

²This metaphor is to say that, if the time complexity of the learner and classifier are known or measurable, then this information can be used to greedily reduce the training error.

the probabilistic equivalent of Adaboost that will serve as reference for the rest of the article. After presenting our main theory in Sections 4 and 5, we discuss further the findings of this study and open directions for future research.

2 Probabilistic weak learner

We consider learning algorithms \mathcal{A} that, given a training dataset $\mathcal{S} = \{(X_1, y_1, D(1)), (X_2, y_2, D(2)), \dots, (X_N, y_N, D(N))\}$, with $D(1) + \dots + D(N) = 1$, return a probabilistic classifier, or oracle, written $h(X)$. For any input X , $h(X)$ is identified with a Bernoulli random variable with parameter $q(+, X)$.

Definition: We say that \mathcal{A} is a *probabilistic weak learner*, if there exists a constant $0 < \varepsilon \leq \frac{1}{2}$ such that, for any dataset \mathcal{S} , the expected error of $h(X)$, is smaller than $\frac{1}{2} - \varepsilon$; that is, one has:

$$\sum_{n=1}^N D(n) q(-y_n, X_n) \leq \frac{1}{2} - \varepsilon,$$

where $q(-y_n, X_n)$ is the probability that $h(X_n)$ takes the value $-y_n$, i.e. that the classifier is wrong.

The constant ε , called the *advantage* or *edge* is unknown and does not need to be known. The probability $q(y_n, X_n)$, also unknown, will be needed. We estimate it by calling repeatedly the weak classifier and calculating the maximum likelihood (ML) or maximum a-posteriori (MAP) estimates: If $h_{1,n}, \dots, h_{R,n}$ are the values returned by R invocations (observations) of $h(X_n)$, then the ML estimate is $\tilde{q}(y, X_n) = |\{n \mid h_{r,n} = y\}| / R$, where $|\cdot|$ is the set cardinal. Assuming that $q(y, X_n)$ is uniformly distributed in $[0, 1]^3$, the MAP is $\hat{q}(y, X_n) = \frac{1}{R+2} (1 + |\{n \mid h_{r,n} = y\}|)$.

3 Adaboost for probabilistic weak learners

We now adapt Adaboost [1] to probabilistic -rather than deterministic- weak learners. Like the original Adaboost, we consider classifiers of the form

$$H(X) = \sum_{t=1}^T \alpha_{t,X} h_t(X),$$

but here, $h_t(X)$ is a Bernoulli random variable (or randomized classifier or oracle), so that $H(X)$ is itself a random variable. Like in Adaboost, we consider domain-partitioned weights (see [1, Sec. 4.1]): we have constants $\alpha_{t,+}$ and $\alpha_{t,-}$ such that $\alpha_{t,X} = \alpha_{t,+}$ if $h_t(X)$ is observed to be $+1$, and that $\alpha_{t,X} = \alpha_{t,-}$ otherwise.

We proceed as in Adaboost, increasing the number T of weak classifier, and not changing a weak classifier once it has been trained. Each random classifier $h_t(X)$ is

³This prior is pessimistic, since $q(y_n, X_n)$ has (unknown) expectation smaller than $\frac{1}{2}$, but the edge ε being unknown, using another prior would not be less hazardous.

obtained by running the weak learner on the training data set $(X_1, y_1), \dots, (X_N, y_N)$, with weights $D_t(n)$, $1 \leq n \leq N$ chosen to emphasize misclassified examples. The weight update rule is:

$$D_{t+1}(n) = D_t(n) (q(t, +, X_n) e^{-\alpha_{t,+} y_n} + q(t, -, X_n) e^{\alpha_{t,-} y_n}) / Z_t, \quad (1)$$

where $Z_t = \sum_{n=1}^N D_t(n) (q(t, +, X_n) e^{-\alpha_{t,+} y_n} + q(t, -, X_n) e^{\alpha_{t,-} y_n})$ normalizes the weights so they sum to one.

With a deterministic weak classifier, one would have $q(\pm, X_n) \in \{-1, +1\}$, resulting in the original Adaboost weight update rule. An additional difference is that the $q(\pm, X_n)$ are unknown. We address this issue in Sec. 3.2 and assume for now that we have estimates $\hat{q}(\pm, X_n)$.

3.1 Boosting property

We now give an upper bound for the expected misclassification error of $H(X)$, and show how to choose the weights $\alpha_{t,\pm}$. This derivation parallels that of [1]: the expected training error is

$$E(\text{Loss}) = \sum_{n=1}^N D(n) E(\mathbb{I}[H(X_n) \neq y_n]) \leq \sum_{n=1}^N D(n) E(e^{-H(X_n) y_n}). \quad (2)$$

where $\mathbb{I}[\cdot]$ is the “indicator function,” being 1 if the bracketed expression is true and zero otherwise.

Since $H(X_n)$ may take at most 2^T possible values, depending on the outputs s_t , $1 \leq t \leq T$, of the T classifiers $h_t(X_n)$, one has:

$$\begin{aligned} E(e^{-H(X_n) y_n}) &= \sum_{s_1, \dots, s_T} \prod_{t=1}^T \hat{q}(t, s_t, X_n) e^{-\alpha_{t,s_t} y_n} \\ &= \sum_{s_1, \dots, s_{T-1}} \prod_{t=1}^{T-1} \hat{q}(t, s_t, X_n) e^{-\alpha_{t,s_t} y_n} \frac{D_{T+1}(n) Z_T}{D_T(n)} \\ &= \sum_{s_1, \dots, s_{T-2}} \prod_{t=1}^{T-2} \hat{q}(t, s_t, X_n) e^{-\alpha_{t,s_t} y_n} \frac{D_T(n) Z_{T-1}}{D_{T-1}(n)} \frac{D_{T+1}(n) Z_T}{D_T(n)} \\ &= \dots \text{etc} \dots \\ &= \frac{D_{T+1}(n)}{D_1(n)} \prod_{t=1}^T Z_t. \end{aligned}$$

Summing over all samples, one gets the familiar expression

$$E(\text{Loss}) \leq \prod_{t=1}^T Z_t.$$

The rest goes as with Adaboost: each Z_t is minimized by setting

$$\alpha_{t,+} = \frac{1}{2} \log \left(\frac{W_t^{++}}{W_t^{+-}} \right) \quad \text{and} \quad \alpha_{t,-} = \frac{1}{2} \log \left(\frac{W_t^{--}}{W_t^{-+}} \right),$$

where $W_t^{ab} = \sum_{n|y_n=b} D_t(n) \hat{q}(t, a, n)$, for any $a \in \{+1, -1\}$, $b \in \{+1, -1\}$. For this choice of $\alpha_{t,\pm}$, one has $Z_t = 2\sqrt{W_t^{++}W_t^{+-}} + 2\sqrt{W_t^{-+}W_t^{--}}$, and one can show that $Z_t \leq \sqrt{1 - 4\varepsilon^2} = \rho$. The expected error of the T -stage boosted probabilistic classifier thus has the same bound as the error of Adaboost:

$$E(\text{Loss}) \leq \left(\sqrt{1 - 4\varepsilon^2}\right)^T \triangleq \rho^T \quad (3)$$

It must be made clear that, in practice, during training, the users only have *estimates* of $q(t, \pm, n)$, so that they reduce *an estimate* of the bound of the expected error.

3.2 Estimation of $q(\pm, X_n)$ during training

In this section, we show how users can, in practice, balance their need for accurate estimates of the $q(t, \pm, n)$ with their eagerness to reduce the estimate on the bound of the expected error (the reader may skip this part in a first reading).

The difference with respect to Adaboost is that, once the classifier has been trained, the user has to estimate the $q(T, \pm, n)$, in order to compute $D_{T+1}(n)$ for the next classifier. The question is thus “how many samples of $h_T(X_n)$ should be taken?” The trivial answer, which we exclude, is to fix some number R of samples and use the corresponding MAP or ML estimates. We exclude for now this approach, to avoid adding an extra parameter R . We propose, instead, two approaches based on the MAP estimator of $q(t, \pm, n)$.

Let us first compare the MAP and ML estimators, to later better explain our preference for the MAP. Both the MAP and ML converge in probability to the true value, so that the corresponding estimators of Z_t also converge in probability to the true value. The MAP and ML differ in that the MAP estimator of $q(T, \pm, n)$ is biased towards $\frac{1}{2}$, and that of Z_T is biased towards 1. More precisely, the expected value of these MAP estimators converge to the their limits *from above*, so that the expected value of successive estimates of Z_T decreases towards the true value. Thus, after sampling $h_T(X_n)$ R times, sampling once more is always expected to decrease the estimate of Z_T .

First approach to estimate $q(T, \pm, n)$: The MAP thus has the advantage of providing a natural stopping time, that of the first observed increase in our estimate of Z_T . The event that Z_T increases has a probability that increases towards $1/2$, so that it will almost always (in the probabilistic sense) happen after a finite time. This strategy can also be used with the ML estimator, but, having a greater variance, it is more likely to result in a spuriously low estimate of Z_T and early stops. On these grounds, the MAP should thus be preferred over the ML.

Second approach: An alternative method involving some look-ahead, and the user’s stopwatch, may be also be considered: having until now trained T classifiers and sampled R times $h_T(X_n)$, $1 \leq n \leq N$, the user has the following options:

- A Train a new classifier $h_{T+1}(X)$, using the current estimate of $q(T, \pm, n)$ in the calculation of $D_{T+1}(n)$. Then sample $h_{T+1}(X_n)$, $1 \leq n \leq N$ once, resulting in a first MAP estimate of Z_{T+1} . As a result, the user decreases the estimated

bound, now $\prod_{t \leq T+1} Z_t$, previously $\prod_{t \leq T} Z_t$, by the factor Z_{T+1} . Also, the user measured, with his or her stopwatch, the elapsed time S_A during training and sampling. The instantaneous bound decrease rate per unit of time is $(Z_{T+1})^{1/S_A}$.

- B** Sample once more $h_T(X_n)$, $1 \leq n \leq N$, producing a new estimate Z'_T . As a result, the user decreases (or increases) the estimated bound by a factor Z'_T/Z_T . Again, with his or her stopwatch, (s)he measured the elapsed time S_B . The instantaneous bound decrease rate is $(Z'_T/Z_T)^{1/S_B}$.

Finally, based on the smallest bound decrease rate, the user decides whether to keep the new classifier $h_{T+1}(X)$ or the new estimate Z'_T .

We have thus proposed two parameterless ways to estimate $q(T, \pm, n)$.

4 Boosting decision tree

Having shown how Adaboost can be transposed to probabilistic weak learners, we now further extend our study to probabilistic decision trees.

Computationally, our proposed classifier is a smoothed binary decision tree. In that model, the output $H(X)$ is a weighed sum of the (random) classifiers on the nodes traversed by an input element X :

$$H(X) = \sum_{t=1}^{T(X)} \alpha_{s(t,X)} h_{s(t-1,X)}(X), \quad (4)$$

where $s(t, X)$ is the index of the t^{th} node reached by input X , $h_{s(t-1,X)}(X) \in \{-1, 1\}$ is the output of the corresponding classifier and $T(X)$ is the depth of the last inner node reached by X before exiting the decision tree. The weight given to $h_{s(t-1,X)}(X)$, $\alpha_{s(t,X)}$ is domain-partitioned, since it depends on the observed value of $h_{s(t-1,X)}(X)$.

Some notation is needed: the index of a node s is a sequence of “+” and “−”, indicating the path to that node. For example, in Fig. 1, right, $s = (+, +, -)$ is the leaf reached by following the “+” edge out of the root node, then the “+” edge out of the (+) node, then the “−” edge out of the (+, +) node. The output of the classifier, when an input exits the decision tree by s , is thus $H(X) = \alpha_+ + \alpha_{++} + \alpha_{++-}$.

For additional convenience, we write \bar{s} the index of the parent of node s ($\bar{s} = (+, +)$ in the previous example) and \dot{s} the last edge followed to reach s (here, $\dot{s} = -$). Thus, one may write $s = (\bar{s}, \dot{s})$. The root node is $s(0, X) = \emptyset$. With this notation, and noting that $H(X)$ only depends on the leaf $l(X)$ reached by X , one has:

$$H(X) = \sum_{\emptyset < s \leq l(X)} \alpha_s \dot{s} \triangleq H_l, \quad (5)$$

where the sum is taken over all nodes s between the leaf l and the root \emptyset (exclusive).

4.1 Decision trees with probabilistic nodes

Like most other decision tree-building algorithms [11, 12], we add nodes one at a time and do not modify previously added nodes. This is the most common way of avoiding the inherent complexity [13] of building decision trees. We do not consider a subsequent pruning step. Unlike other decision trees, and like in Adaboost, each node is trained on the whole dataset.

However, we modulate the weights of the examples, not only based on whether they are misclassified (as in Adaboost), but also based on their probability of reaching the node. After having trained $h_s(X)$ with weights $D_s(n)$, $1 \leq n \leq N$, the weights for training the children nodes $s+$ and $s-$ are:

$$D_{s+}(n) = \frac{D_s(n)}{Z_{s+}} q(s+, X_n) e^{-\alpha_{s+} y_n} \text{ and } D_{s-}(n) = \frac{D_s(n)}{Z_{s-}} q(s-, X_n) e^{-\alpha_{s-} y_n}. \quad (6)$$

In this expression, $Z_{sa} = \sum_{n=1}^N D_s(n) q(sa, X_n) e^{-\alpha_{sa} y_n}$, for $a \in \{+, -\}$, are normalizing constants, and $q(s+, X_n) \in [0, 1]$ is the (unknown) parameter of the Bernoulli random variable $h_s(X)$. Like in Sec. 3, we use estimates $\hat{q}(s+, X_n)$ in place of the true values.

4.2 Bound on the expected error

We now bound the error of the boosting tree algorithm and specify the weights α_s and the choice of the trained node at each step.

Using again the exponential error inequality $\text{Loss}(H(X), y) \leq e^{-H(X)y}$, Eq. (2), the expected misclassification error for a training example X_n is upper-bounded by

$$E\left(e^{-H(X_n)y}\right) = \sum_{l: \text{leaf } H} p(l, X_n) e^{-H_l y_n}, \quad (7)$$

where $p(l, X)$ is the probability of an input X reaching the leaf l . More generally, assuming independence of the outputs of classifiers at each node, the probability that X reaches a node $s = (s_1, s_2, \dots, s_T)$ is

$$p(s, X) = q(s_1, X) \cdot q(s_1 s_2, X) \cdot \dots \cdot q(s_1 \dots s_D, X) = \prod_{r \leq s} q(r, X),$$

where the product is taken for all nodes between the root and s .

The error bound is thus

$$\begin{aligned} E\left(e^{-H(X_n)y}\right) &= \sum_{l: \text{leaf } H} \left(\prod_{s \leq l} q(s, X_n) \right) e^{-H_l y_n} \\ &= \sum_{l: \text{leaf } H} \prod_{s \leq l} (q(s, X_n) e^{-\alpha_s y_n}) \\ &= \sum_{l: \text{leaf } H} \prod_{s \leq l} \left(\frac{D_s(n)}{Z_s} \right) \\ &= \sum_{l: \text{leaf } H} D_l(n) \prod_{s \leq l} Z_s \end{aligned}$$

Summing over all examples X_n , $1 \leq n \leq N$ and replacing in Eq. (2) yields the bound:

$$E(\text{Loss}) \leq \sum_{l: \text{leaf } H} \prod_{s \leq l} Z_s \quad (8)$$

Like above, each Z_s is minimized by setting

$$\alpha_s = \frac{1}{2} \log \left(\frac{W_s^{\dot{s}\dot{s}}}{W_s^{\dot{s}\neg\dot{s}}} \right), \text{ where } W_s^{ab} = \sum_{n|y_n=b} D_s(n) q(a, n), \quad a, b \in \{+, -\},$$

and \neg is the negation operator. For these values of α_s , each Z_s takes the value

$$Z_s = 2\sqrt{W_s^{\dot{s}+} W_s^{\dot{s}-}}.$$

This bound can also be found, in slightly different contexts, in our previous work [14] and in our unpublished manuscript [15]. In the present paper, we additionally study how this bound evolves with the size of the tree.

Expected error bound as a function of the tree size

We now describe the evolution of the bound (8) when the tree is grown by a greedy bound-reducing algorithm.

As previously, we may show that $Z_{s+} + Z_{s-} \leq \sqrt{1 - 4\varepsilon^2} = \rho$, owing to the probabilistic weak learner hypothesis.

We now proceed recursively. After training and incorporating T nodes, the expected error bound is $C(T) = \sum_{l: \text{leaf } H} \prod_{s \leq l} Z_s$. At this point, the tree has $T + 1$ leaves, so that one leaf l at least has an error not less than $C(T) / (T + 1)$. After training a probabilistic weak classifier at l , the new error bound is

$$\begin{aligned} C(T+1) &= C(T) - \prod_{s \leq l} Z_s + \prod_{s \leq l-} Z_s + \prod_{s \leq l-} Z_s \\ &= C(T) + \left(\prod_{s \leq l} Z_s \right) (-1 + Z_{l+} + Z_{l-}) \\ &\leq C(T) + \left(\frac{C(T)}{T+1} \right) (-1 + \rho) \\ &= C(T) \left(\frac{T+\rho}{T+1} \right). \end{aligned}$$

Since $C(0) = 1$, we have the general relation

$$C(T) \leq \prod_{t=0}^T \frac{t+\rho}{t+1} = \frac{1}{TB(T, \rho)} \triangleq F(T, \rho) \simeq \frac{T^{\rho-1}}{\Gamma(\rho)}, \quad (9)$$

where $B(T, \rho)$ is the beta function and $\Gamma(\rho)$ is the Gamma function. The rightmost term is the asymptotic approximation for large T ; it is coherent with the bound of d[10, Eq. 6].

This bound is interesting in more than one respect:




Figure 2: Left: Bound of boosted decision tree (full curve, highest), Eq. (9), of probabilistic Adaboost (dotted, lowest), Eq. (3), and of matryoshka (dashed, middle). From top to bottom, $\rho \in \{\frac{31}{32}, \frac{7}{8}, \frac{3}{4}, \frac{1}{2}, \frac{1}{4}\}$, i.e. $\varepsilon \in \{0.12, 0.24, 0.33, 0.43, 0.46\}$. Right: Bound of boosted tree of simple trees, given by Eq. (10).

- It appears that it cannot be very much improved, in the following sense: consider a probabilistic learner with error $1/2 - \varepsilon$, independently of the weights $D(n)$ with which it is trained. This learner verifies the probabilistic weak learner hypothesis. Now, for both the probabilistic Adaboost and for a (balanced) decision tree, $H(X)$ is a binomial random variable with parameters $(1/2 - \varepsilon)$, and the number of weak parameters traversed by X . This second parameter is T for Adaboost and $\log_2(T)$ for the decision tree. It is clear, then, that the decision tree requires exponentially more weak classifiers than the probabilistic Adaboost.
- This bound is especially bad for very weak classifiers ($\rho \simeq 1$). The full curves in Figure 2, left, plot the bound $F(T, \rho)$ for $\rho = 31/32, 7/8, 3/4, 1/2$ and $1/4$. For comparison, the expected error bound of Adaboost, ρ^T , plotted alongside, is much lower, especially for $\rho = 31/32$.
- This bound calls the attention of designers of decision trees tempted to pass all the training dataset along all branches: if the weak classifier is very weak, the number of needed weak classifiers may grow very much. With stronger classifiers, the boosting tree algorithm may be more practical.

5 Matryoshka decision trees

Based on the conclusion of the previous section -that stronger classifiers yield better boosted decision trees, we now address the question of obtaining sufficiently strong classifiers. The first step in this direction (Section 5.1) is to explore the idea of putting a boosted tree at each node. We will see that there is an advantage in doing so. It will then be natural, in Section 5.2, to build trees of trees of trees of ... of weak classifiers, that is, a matryoshka of decision trees.

5.1 Bound for a tree of trees

In this section, we study the error bounds obtainable by a decision tree built using the method of Section 4, but where the nodes are themselves trees built according to that

same method. We place ourselves in the situation of having the resource to train a fixed number T of weak classifiers, and our objective is to minimize the bound on the expected error.

In this context, it is natural to study the bound obtainable by assembling T_2 sub-trees of fixed size $T_1 = T/T_2$. By Eq. (9), the error bound for the sub-trees is $F(T_1, \rho) = (T_1 B(T_1, \rho))^{-1}$, and that of the outer tree is

$$F\left(\frac{T}{T_1}, F(T_1, \rho)\right). \quad (10)$$

Figure 2, right, plots this bound plotted against T_1 . The curves show that, for $T_1 = 1$ and $T_1 = T$, the bound is the same as $F(T, \rho)$, i.e. that of a not-nested decision tree. More interestingly, for intermediate values of T_1 , the bound of Eq. (10) is always lower than $F(T, \rho)$. In particular, the minimum is always near $T_1 = \sqrt{T}$.

Given these encouraging results, we are naturally tempted to substitute the sub-trees (of size T_1) by T'_1 sub-trees of sub-sub-trees of size T''_1 , for some T'_1, T''_1 s.t. $T'_1 T''_1 = T_1$. The same idea can also be applied to the outer tree.

5.2 Bound for a tree of trees ... of trees of weak classifiers

More generally, we are tempted to determine the bounds reachable by trees of trees of ... of trees of weak classifiers. For some L and T_1, T_2, \dots, T_L s.t. $T_1 T_2 \dots T_L = T$, the bound is easily shown to be:

$$F(T_L, F(T_{L-1}, \dots F(T_1, \rho))).$$

Finding analytically the optimal combination of T_i , $1 \leq i \leq L$ may not be easy. But, guided by the observation that, for $L = 2$, the optimal choice seems to be near $T_1 = T_2 = \sqrt{T}$, we naturally consider the case $T_1 = T_2 = \dots = T_L = T^{1/L}$. In this case, the bound is

$$F\left(T^{\frac{1}{L}}, F\left(T^{\frac{1}{L}}, \dots F\left(T^{\frac{1}{L}}, \rho\right)\right)\right). \quad (11)$$

The black graph in Figure 3 plots this value against the nesting level L , with the original bound $F(T, \rho)$ (top) for comparison. This figure clearly shows that deeper nesting levels improve the bound. In fact, Eq. (11) continues to decrease for $L > \log_2 T$, i.e. when the trees each have less than two nodes.

This (strange) effect is due to the fact that $F(T, \rho)$ is defined for any positive real T . Since the number of nodes is in an integer, there are no practical repercussions.

However, these curves clearly indicate that smaller sub-trees yield better bounds. This suggests building the smallest possible trees, with just two nodes, each node a tree with two nodes, etc, until the last level, consisting of trees with two weak classifiers.

5.3 Bound for 2-matryoshka

We now derive the expected error bound for the “2-matryoshka” tree, having exactly two nodes, at all nesting levels, having precisely two nodes. We thus need to assume that $T = 2^L$ is a power of two.




Figure 3: Black curve: bound on error of matryoshka decision trees at various levels of nesting. The sub-tree sizes are $T^{1/(\text{nesting level})}$. Light-colored curves near the black curve are for trees w/ integer number of nodes. . The topmost line marks the error bound of the (not nested) decision tree. At left, $T = 1024$, and $T = 65536$ at right. These curves are for $\rho = \frac{31}{32}$ i.e. $\varepsilon \simeq 0.12$.

We call $M_2(T, \rho) = F(2, F(2, \dots F(2, \rho)))$ the bound for this tree (there are L nested parentheses). Recalling from Eq. (9) that $F(2, \rho) = \rho^{\frac{1+\rho}{2}} = \frac{1}{2}\rho + \frac{1}{2}\rho^2$, one writes $M_2(T, \rho)$ as a polynomial of degree T .

Figure 2, left, shows the graph of $M_2(T, \rho)$, in dashed lines. This figure shows that the 2-matryoshka tree has a much stronger boosting ability than the plain boosting tree, and this is the main result of this paper.

5.4 Building a matryoshka

The algorithm for the 2-matryoshka would thus be: train a two-leaf tree (stage “b”, in Fig. 1), and collect the leaves into a single node (“c”). Train a two-leaf sub-tree on one of the branches, collect its leaves in a single node (“e”). Collect the leaves once more (“f”) etc. If all weak classifiers have the same edge ε , then this approach is the most appropriate.

In practice, the classifiers will not have the same edge and a greedy -with respect to number of nodes or physical training time- bound-decreasing approach could be considered. Each time a classifier is added to the tree, we will consider each sub-tree containing that node, starting from the top. For each sub-tree, we compare the instantaneous bound decrease rate⁴ of the sub-tree at T ,

$$\dot{C}_{Simple} \simeq (C(T+1) - C(T-1)) / 2, \quad (12)$$

($C(T)$ being computed on the sub-tree only), with that of a tree having such a sub-tree at each node,

$$\dot{C}_{Matryoshka} = \frac{\partial}{\partial t} F\left(\frac{t}{T}, C(T)\right) (t = T). \quad (13)$$

If the later is smaller, then the leaves of the sub-tree are collected into a single node.

We now give the detail of computing Eq. (13). Using the relation $\frac{\partial}{\partial x} B(x, y) = B(x, y) (\psi(x) - \psi(x+y))$, where ψ is the digamma function, $\psi(x) = \frac{\partial}{\partial x} (\log(\Gamma(x)))$,

⁴Here, we consider the decrease rate per added node, but the decrease rate per unit of training time could be used too.

one gets

$$\begin{aligned} F'_T(T, \rho) &= -F(T, \rho) \left(\frac{1}{T} + \psi(T) - \psi(T + \rho) \right) \text{ and} \\ F'_\rho(T, \rho) &= -F(T, \rho) (\psi(\rho) - \psi(T + \rho)). \end{aligned}$$

The first line above then gives

$$\frac{\partial}{\partial t} F\left(\frac{t}{T}, C(T)\right) (t = T) = \frac{C(T)}{T} \left(\gamma + \psi(C(T)) + \frac{1}{C(T)} - 1 \right), \quad (14)$$

where $\gamma = -\psi(1) \simeq 0.5772$ is Euler's constant.

One can check that, for $T = 1$, $\dot{C}_{Simple} = \dot{C}_{Matryoshka}$ and that, if $C(T) = F(T, \rho)$, i.e. if the bound (9) is tight, then $\dot{C}_{Simple} = \dot{C}_{Matryoshka}$ for all $T > 1$.

6 Discussion and conclusions

We have developed in this paper a theory of probabilistic boosting, aimed at decision trees. We proposed a boosting tree algorithm and a theoretically superior matryoshka decision tree algorithm. These algorithms are essentially parameter-free, owing to the principle of choosing whichever training action most reduces the expected training error bound, and to a judicious choice of possible training actions.

We showed bounds on the expected training error of the algorithms, one of them discouraging, the other, encouraging. The bounds for simple trees and for trees of trees are coherent with our early experiments.

Future developments include an analysis of the effect of approximating the node branching probabilities $q(s, X_n)$ during training and experimental evaluation of the matryoshka.

On a more general level, we believe that the high bound for boosting trees indicates that the probabilistic weak learner hypothesis is inadequate. This hypothesis, directly adapted from the theory of boosting, does not take into account the fact that real-world classifiers usually have a lower training error on smaller training sets. Our intuition is thus that the entropy of the training weights, $D(n)$, should be taken into account in future work.

References

- [1] R. E. Schapire and Y. Singer. Improved boosting algorithms using confidence-rated predictions. *Machine Learning*, 37(3):297–336, 1999.
- [2] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.
- [3] L. Breiman. Random forests. *Machine Learning*, 45:5–32, 2001.
- [4] Z. Tu. Probabilistic boosting-tree: Learning discriminative models for classification, recognition, and clustering. In *proc. ICCV*, 2005.
- [5] P. Viola and M. Jones. Robust real-time object detection. In *proc. ICCV workshop on statistical and computational theories of vision*, 2001.

- [6] B. McCane and K. Novins. On training cascade face detectors. In *Image and Vision Computing New Zealand*, 2003.
- [7] H. Luo. Optimization design of cascaded classifiers. In *proc. CVPR*, 2005.
- [8] R. L. P. Chang and T. Pavlidis. Fuzzy decision tree algorithms. *IEEE Trans. Systems, Man, and Cybernetics*, 7(1):28–35, 1977.
- [9] J. R. Quinlan. *Probabilistic decision trees*, in *Machine Learning: An Artificial Intelligence Approach*, volume 3, chapter 5, pages 140–152. Morgan Kaufmann, 1990.
- [10] M. Kearns and Y. Mansour. On the boosting ability of top-down decision tree learning algorithms. *J. of Computer and Systems Sciences*, 58(1):109–128, 1999.
- [11] J. R. Quinlan. *C4.5 : Programs for machine learning*. Morgan Kauffann, 1993.
- [12] T. M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- [13] L. Hyafil and R.L. Rivest. Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5(1):15–17, 1976.
- [14] E. Grossmann. AdaTree : boosting a weak classifier into a decision tree. In *Workshop on Learning in Computer Vision and Pattern Recognition, CVPR*, 2004.
- [15] E. Grossmann. Adatree 2 : Boosting to build decision trees, or Improving Adatree with soft splitting rules. unpublished work done at the Center for Visualisation and Virtual Environments, University of Kentucky, 2004.